Oxford Journal of Legal Studies 2022, Vol. 42, No. 3 pp. 869–892
https://doi.org/10.1093/ojls/gqac004

What the Tortoise Says about Statutory Interpretation: The Semantic Canons of Construction Do Not Tip the Balance

Amin Ebrahimi Afrouzi*

Abstract—Karl Llewellyn's critique of the canons of statutory interpretation led to a decline in their use for several decades. His critique, however, faced sustained resistance from some corners of the academy and the judiciary. Although this resistance has had only a selective uptake, it animated a gradual revival of the canons and brought the state of scholarship to an impasse that is for the most part partisan. In this article, I examine the semantic canons from a deeper level and argue that a universal assumption about them is false. Said assumption is that, although not dispositive, the semantic canons at least offer some reasons in favour of or against a candidate interpretation. Inclinations to rely on the semantic canons are also based on this assumption, although it is an assumption that the critics of the canons also share. I argue that this assumption is false because the semantic canons are a class of rules that are by nature not reason-giving. This provides a new ground against giving the semantic canons deliberative weight in questions of statutory interpretation.

Keywords: statutory interpretation, legislation, textualism, legal philosophy, normativity, legal reasoning

1. Introduction

The semantic canons of construction, also collectively known as linguistic, textual or syntactic canons of construction, are a collection of maxims, many with Latin names, that are routinely invoked in questions of statutory interpretation. For example, the canon *ejusdem generis* (literally, 'of the same type') maintains that when a general term appears at the end of a list of specific examples, its meaning is limited by the shared characteristics of the examples in the list. A classic instance of this canon is Justice Holmes's reasoning in *McBoyle*, holding an airplane not to be a motor vehicle within the meaning of a statute reading: 'the term "motor vehicle" shall include an automobile, automobile truck, automobile

* PhD Candidate, University of California at Berkeley, School of Law, Program in Jurisprudence and Social Policy. Email: am1n@berkeley.edu. I am indebted to Christopher Kutz and Bertrall Ross for reading several earlier drafts and giving me extensive feedback on each; to Barry Stroud, John Searle, James S Dillon, Matthias Brinkmann, Melissa Fusco, James Mathura, Raha Ahmadian, David Singh Grewal, Jeffrey Kaplan, Haley Anderson and an anonymous referee for generous feedback on earlier drafts; and to Clara Choi for copy-editing and cite-checking. All faults remain mine.

wagon, motor cycle, or any other self-propelled vehicle not designed for running on rails'.[1] Though Justice Holmes did not refer to the canon by its Latin name, he invoked it by reasoning that because all the vehicles in the list run on land, the term 'motor vehicle' must also be limited to those vehicles that run on land.[2]

A similar canon, called *noscitur a sociis* (literally, 'known from associates'), maintains that a word's meaning can be discerned by way of reference to its surrounding words. Thus, a different meaning for the word 'banks' would be discerned if it appears in a context regarding financial institutions rather than rivers.[3] Typically, this canon is said to apply where a statutory term can have both a broad and a narrow meaning, 'and the surrounding terms may offer some clue as to which meaning was intended'.[4]

These and other semantic canons are often invoked as reasons favouring a certain statutory interpretation over alternatives. Although they are often not regarded as dispositive, some judicial decisions are made by relying entirely on a semantic canon. Unlike substantive canons of constructions, such as the rule of lenity,[5] the collective appeal of the semantic canons lies in their presumed substantive neutrality. While the majority of questions before the federal courts are those of statutory interpretation[6] and the use of the semantic canons has been on the rise,[7] strong voices from the academy advocate for a formal system of *stare decisis* to make the canons binding on lower courts.[8] Against this background, I

[1] *McBoyle v United States* 283 US 25, 26 (1931).

[2] In fact, Nolo's *Plain-English Law Dictionary* defines the canon in terms of this case: 'Ejusdem Generis' <www.nolo./dictionary/ejusdem-generis-term.html> accessed 31 January 2019.

[3] See Lawrence Lessig, 'Understanding Changed Readings: Fidelity and Theory' (1995) 47 Stan L Rev 395, 407–8.

[4] John F Manning and Matthew C Stephenson, *Legislation and Regulation* (2nd edn, Foundation Press 2010) 226 (discussing *Gustafson* as the example of such a case, where the majority construed 'communication' as only referring to a message that is widely disseminated, a narrow reading, as opposed to referring to just anything that conveys information from one party to another, a broad reading).

[5] Substantive canons are rules that codify substantive preferences about what kinds of answers a court ought to prefer in determining statutory meaning. For further definitions of substantive canons, see eg ibid 202, which defines them as 'judicial presumption[s] … in favor of or against a particular substantive outcome'; William N Eskridge, Jr and Philip P Frickey, *Cases and Materials on Legislation: Statutes and the Creation of Public Policy* (2nd edn, West Academic Publishing 1995) 634, describing them as 'essentially presumptions about statutory meaning based upon substantive principles or policies drawn from the common law, other statutes, or the Constitution'.

[6] See Ruth Bader Ginsburg, 'Informing the Public about the US Supreme Court's Work' (1998) 29 Loy U Chi LJ 275, 282.

[7] See generally Aaron-Andrew P Bruhl, 'Communicating the Canons: How Lower Courts React When the Supreme Court Changes the Rules of Statutory Interpretation' (2015) 100 Minn L Rev 481 (especially part II, showing that lower federal courts followed suit as the use of semantic canons in statutory interpretation by the Supreme Court gained momentum).

[8] Jordan Wilder Connors, 'Treating Like Subdecisions Alike: The Scope of Stare Decisis as Applied to Judicial Methodology' (2008) 108 Colum L Rev 681, 708–14; Sydney Foster, 'Should Courts Give Stare Decisis Effect to Statutory Interpretation Methodology?' (2009) 96 Geo LJ 1863, 1866–7. See also Abbe R Gluck, 'Intersystemic Statutory Interpretation: Methodology as "Law" and the Erie Doctrine' (2011) 120 Yale LJ 1898, 1909–18 (discussing whether interpretive methodology is binding law); Nicholas Quinn Rosenkranz, 'Federal Rules of Statutory Interpretation' (2002) 115 Harv L Rev 2085, 2145–56 (arguing for the codification of binding interpretive rules). In a more recent article, however, Gluck points out that 'Case law, empirical work, and judicial writings all confirm that most judges have a visceral, highly negative reaction' to the formalisation of the methods of statutory interpretation whether it be done by the Supreme Court or Congress: Abbe R Gluck, 'Congress, Statutory Interpretation, and the Failure of Formalism: The CBO Canon and Other Ways That Courts Can Improve on What They Are Already Trying to Do' (2017) 84 U Chi L Rev 177, 184. See also Bruhl (n 7) 487–91 (stating that this may be *de facto* happening, even in the absence of a formal structure).

argue that the semantic canons of construction belong to a class of rules that are by nature not reason-giving, or *normative*. Accordingly, giving them any weight in questions of statutory interpretation is a mistake.

General suspicion about the canons of construction is, of course, nothing new. Karl Llewellyn's critique of the 'dueling canons', for example, cast a shadow on their use for many decades.[9] More recently, Richard Posner has called the canons 'Vacuous and inconsistent' and 'just plain wrong'.[10] These critiques, however, face sustained resistance.[11] For example, Justice Scalia, though qualifying them as indecisive, famously wrote that relying on the semantic canons is 'so commonsensical that, were the canons not couched in Latin, you would find it hard to believe that anyone could criticize them'.[12] Focusing only on the semantic canons, I enter this debate from a deeper theoretical level and question one of its underlying premises. Said premise is that the semantic canons, even if not dispositive, still offer some reasons in favour for or against a given interpretation of statutory language. This presupposes that the semantic canons are *reason-giving* rules.

A rule is reason-giving, or normative, when we have reasons in light of the rule that we would not have without it.[13] But not all rules are reason-giving in this sense.[14] Non-reason-giving rules bear no weight in deliberation. I will argue that the semantic canons belong to this latter class of rules. This allows me to articulate exactly why it is a mistake to rely on the semantic canons in statutory interpretation. It does not, however, put me in the company of existing critics of the canons, for, as my analysis reveals, these critics operate under the same misconception as the canons' proponents.

Karl Llewellyn, who was the most notable among such critics, also advocated against using the semantic canons altogether.[15] His grounds for doing so, however, were both mistaken and different from mine. He wrote that 'the construction contended for must be sold, essentially, by means other than the use of the

[9] Karl N Llewellyn, 'Remarks on the Theory of Appellate Decision and the Rules or Canons about How Statutes Are to Be Construed' (1950) 3 Vand L Rev 395, 401. See also Manning and Stephenson (n 4) (discussing the dooming impact of Llewellyn's critique on the canons).

[10] Richard A Posner, 'Statutory Interpretation—in the Classroom and in the Courtroom' (1983) 50 U Chi L Rev 800, 806, 816.

[11] eg Antonin Scalia, 'Common-Law Courts in a Civil-Law System: The Role of United States Federal Courts in Interpreting the Constitution and Laws' in Amy Gutmann (ed), *A Matter of Interpretation: Federal Courts and the Law* (Princeton UP 1997) 3, 27; Adam Schlusselberg and Michael BW Sinclair, '"Only a Sith Thinks Like That": Llewellyn's "Dueling Canons," Twenty Five to Twenty Eight' (2010) NYLS Legal Studies Research paper No 10/11 #3.

[12] Scalia (n 11) 26.

[13] See the distinction between 'formal' and 'robust' norms in Scott Shapiro and David Plunkett, 'Law, Morality, and Everything Else' (2017) 128 Ethics 37. For predecessors of this distinction, see Tristram Mc-Pherson, 'Against Quietist Normative Realism' (2011) 154 Philosophical Studies 223; Tristram McPherson, 'Authoritatively Normative Concepts' in R Shafer-Landau (ed), *Oxford Studies in Metaethics No 13* (OUP 2017). See also the thin definition of 'norm' given in Scott Shapiro, *Legality* (Harvard UP 2013). For connected discussion, see David Copp, 'Moral Naturalism and Three Grades of Normativity' in Peter Schaber (ed), *Normativity and Naturalism* (Ontos-Verlag 2005) 7–46. See also Derek Parfit, *On What Matters*, vol 1 (OUP 2011), who contrasts normativity in the 'rule-implying' sense with normativity in the 'reason-implying' sense.

[14] Shapiro and Plunkett (n 13) call these 'formal rules'.

[15] Llewellyn's argument was also targeted against using the substantive canons, but this is of no interest here because my focus is on the semantic canons.

canon', but only due to scepticism about whether one could reasonably *choose* between what seemed to be equally forceful applicable canons.[16] Llewellyn's conclusion is based on his finding that 'there are two opposing canons on almost every point', which suggests a belief to the existence of two reasons in opposite directions.[17] Moreover, the suggestion here is not that the two reasons cancel each other out, but that one of them outweighs the other, though with no criteria discernible by us. He concluded, therefore, that the choice between the two seemingly applicable canons must be made arbitrarily.[18] This conclusion led Llewellyn and other so-called 'legal realists' to condemn the canons as disingenuous tools by which judges couch their political views in ostensibly objective terms.[19] But this did not lead them to view the canons as *void* of reasons, which is what I argue they are.[20]

Much contemporary scholarship has Llewellyn's critique as its backdrop. Scalia (and others) directly responded to Llewellyn's critique by claiming that he misrepresents narrow exceptions to broader canons as 'counter-canons', arguing that this merely shows that semantic canons, like all rules, are not without exceptions.[21] Macey and Miller respond to Llewellyn's scepticism by writing that 'it seems plausible that judges can select among [contradictory] canons in a sensible and coherent fashion even in the absence of known rules to guide them'.[22] Eskridge and Frickey characterise the canons as rules that are in need of further rules before they can be relied on.[23] Brudney, who maintains that questions of statutory interpretation must be resolved based on reliance on the right canon, still cautions that such further rules may also be arbitrary, writing that 'The Court's failure to develop any interpretive rubric for prioritizing or ordering its reliance on different canons may well stem from an implicit understanding that such a creation would be both arbitrary and unproductive'.[24] Manning and Stephenson worry whether further rules about choosing between canons would lead to an infinite regress, writing:

[16] Llewellyn, 'Remarks on the Theory of Appellate Decision' (n 9) 401.
[17] ibid.
[18] ibid.
[19] See generally ibid. See also *Cont'l Cas Co v Pittsburgh Corning Corp* 917 F.2d 297, 300 (7th Cir 1990) (calling the canons 'figleaves for decisions reached on other grounds').
[20] For further discussion of this, see Posner (n 10) 806, where, though accepting Llewellyn's critique wholesale, Posner goes on to reject many substantive canons and semantic canons on substantive or empirical grounds. Specifically, he rejects the semantic canons he considers by arguing that they are 'based on the assumption of legislative omniscience', which assumption he deems implausible. Posner (n 10) 812–13. Thus, he considers the semantic canons are not categorically void of reasons but, as he specifically states, 'wrong'. Posner (n 10) 806.
[21] Scalia (n 11) 27. Continuing this line of argument, Schlusselberg and Sinclair more recently wrote: 'Llewellyn's technique in creating the [dueling canons] list was in most cases to take a familiar and acceptable statement about interpretation and split the exception from the main statement, making one the thrust and the other, usually the exception, the parry': Schlusselberg and Sinclair (n 11) 38. It is not clear how this line of defence does not merely displace rather than resolve the problem raised in Llewellyn's critique, but I leave this matter aside.
[22] Jonathan R Macey and Geoffrey P Miller, 'The Canons of Statutory Construction and Judicial Preferences' (1992) 45 Vand L Rev 647, 650–1.
[23] See William N Eskridge, Jr and Philip P Frickey, 'Forward: Law as Equilibrium' (1994) 108 Harv L Rev 26, 66.
[24] James J Brudney, 'Canon Shortfalls and the Virtue of Political Branch Interpretive Assets' (2010) 98 CLR 1199, 1231.

> If we need a second tier of rules of thumb to understand how to apply the first tier of rules of thumb (the canons), then will we eventually need a third tier to understand how to apply the second tier, and so on?[25]

All these views share a feature with Llewellyn and one another: they regard the canons as reasons in favour of or against candidate interpretations of statutory language, albeit reasons that are hard to weigh.

My analysis suggests that the problem is not about succeeding or failing to find (or creating) non-arbitrary criteria in making a choice that must be made. Rather, considering a question of statutory interpretation as a choice between competing canons is already mistaken, for, as I shall argue, none of the supposedly competing canons bear any reasons for or against a candidate interpretation of statutory language.

The intuition that the semantic canons are not normative can already be gleaned from Eskridge and Frickey's observation suggesting the lameness of the canons by themselves, as well as from Manning and Stephenson's caution that higher-level rules would be equally lame. However, this is not a matter that can be settled by intuitions. My aim therefore is to clarify why the semantic canons are indeed non-normative rules. In pursuing this aim, I draw from and apply two important philosophical tools: John Rawls's notion of summary rules and what is known as the Lewis Carroll paradox.

The semantic canons are almost universally described as maxims of interpretation or rules of thumb, or in similar terms.[26] Yet, the study of the canons in the capacity or character of *rules* is practically non-existent. Meanwhile, the philosophical literature about the nature of rules and their normativity is vast. I address this gap and consider canons as a species of rules and examine whether or how they can be reason-giving. Thus, its conclusions aside, this article aims to enhance our understanding of the canons in ways that have not yet been explored in the study of statutory interpretation.

In the simplest terms, I argue that by their nature, the semantic canons do not weigh in on the balance of reasons in questions of statutory interpretation. As will become clear, this also sets them wholly apart from substantive canons and other issues of statutory interpretation.[27] Thus, I shall proceed without recourse to broader matters of policy, interpretation or theory.[28]

[25] Manning and Stephenson (n 4) 214.

[26] See ibid 204 (defining substantive canons as the 'rules of thumb for decoding legal language').

[27] In fact, the confusions about the semantic canons may be partly arising from the fact that most scholars treat the semantic and substantive canons together.

[28] Thus, without relying on anything like what Jerry Mashaw has in mind, when writing that 'Any theory of statutory interpretation is at base a theory about constitutional law'. Jerry Mashaw, 'As if Republican Interpretation' (1988) 97 Yale LJ 1685, 1686. There are larger problems in statutory interpretation that will not be of concern here. I shall not defend any view on how one interpretation should be selected over another when conflicting interpretations are possible, nor on what basis one interpretive path may be chosen when different possible paths lead to conflicting outcomes. All that would be of concern is whether the semantic canons should play any role in resolving questions of statutory interpretation, which I shall answer in the negative.

In section 2, I introduce John Rawls's notion of 'summary rules'[29] and argue that these rules are by nature not reason-giving. I then narrow in on how deliberation goes astray when summary rules are mistaken for reasons by discussing a handful of thought experiments and court cases famous for their reliance on the semantic canons. In section 3, I draw on the Lewis Carroll paradox to argue that the semantic canons are indeed summary rules in the sense specified. Once this is established, in section 4 I extend what was said in section 2 about summary rules to the semantic canons of constructions; specifically, that they do not tip the balance of reasons in questions of statutory interpretation. Sections 2–4 cover the positive argument of this article. Sections 5 and 6 respond to some possible objections.

2. *Summary Rules*

In his 'Two Concepts of Rules', John Rawls points out that 'rules of thumb' or 'maxims' arise from patterns in recurrent cases that are similarly decided by the direct examination of all relevant factors.[30] He calls such rules 'summary rules' because they are 'summaries of past decisions' and, though not apparent in their content, merely report that 'cases of a certain sort have been found on *other* grounds to be properly decided in a certain way'.[31] The key point to understand about these rules is that they have no force in establishing how 'a case of a certain sort' *ought* to be decided. Rather, they report a pattern in the presence of independent grounds that tend to make the deciding of those sorts of cases proper, in the way recounted by them.[32]

Rawls's analysis is dense, so I will unpack it with an example. Suppose that Sally is about to take on an important business venture, but she is unsure whether to embark on it alone or to take Joe as her partner. She remembers the following rule of thumb: 'it takes two to tango'. This rule tracks a pattern in the fact that people often have complementary skill sets. The rule reports that in many situations, partnering with another person will lead one to more success.

This may or may not be the case in the instance of concern; Joe may not be a man of any useful skill. If it turns out in this instance that Joe indeed has useful and complementary skills to Sally, what *makes* partnering with him in this instance appropriate is not the recounted rule, but the underlying fact a pattern

[29] John Rawls, 'Two Concepts of Rules' (1955) 64 The Philosophical Review 3, 19. Rawls's account in this article is embedded in the context of a defence of utilitarianism against some famous critiques. Because this broader aim of the article is not relevant to the discussion here, I shall discuss Rawls's analysis in abstract form throughout. Though incisive and illuminating, Rawls's analysis of the nature of rules in this early article remains underappreciated, partly perhaps due to his change of heart and the overshadowing of his earlier work by his later and much more influential career as a *critic* of utilitarianism.

[30] Rawls (n 29) 23: 'It is doubtful that anything to which the summary conception did apply would be called a *rule*. Arguing as if one regarded rules in this way is a mistake one makes while doing philosophy' (original emphasis).

[31] ibid 19 (original emphasis).

[32] See ibid 22–4. As will be clear in section 3, the issue at hand is different from the problem of induction.

of which the rule tracks, namely the pattern that people often have complementary skills. That fact is a *reason* for Sally to partner with Joe.[33] The rule is not.

Thus, the rule has no force in establishing that the right thing to do, when embarking on an important project, is to partner with another person. Rather, the rule reports that there is a pattern of independent reasons that renders doing so in many instances the right thing to do. The reason in this example is independent from the rule because the fact that people often have complementary skills would not have been any more or less true whether such a rule existed or not.[34]

It may seem that Sally would have reason to follow the rule considered, because it is an apt generalisation. However, several issues must be noted.

First, if Sally is able to discern or evaluate the independent reasons present in every instance of making such a decision, she will have no need for the rule or for the underlying pattern of facts that it invokes. And if this were generally true of persons, summary rules would never have existed. As Rawls notes, a society of experts 'would be a society without [summary] rules in which each [expert] person' performs the relevant analysis on each instance directly, and without error.[35] In this context, 'expert' only means someone who is capable of discerning and evaluating the independent reasons present in each instance with a lower margin of error than the generalisation that can be made based on the pattern reflected in the rule. Such experts have no use for summary rules.[36]

This brings us to the second matter: suppose that Sally is not such an expert, that is, the chances for Sally to make a mistake are higher than the error margin of the pattern underlying the rule. It is only in these cases that relying on the rule would make sense.[37] In fact, the only reason summary rules exist is

[33] In this article, when speaking of reasons, I exclusively mean to denote 'normative reasons', which are reasons that justify a certain decision or action. To use TM Scanlon's definition, a normative reason is 'a consideration that counts in favour of' someone's acting in a certain way: TM Scanlon, *What We Owe to Each Other* (Harvard UP 1998) 17–77. See generally TM Scanlon, 'Reasons: A Puzzling Duality?' in R Jay Wallace, Philip Pettit, Samuel Scheffler and Michael Smith (eds), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz* (Clarendon Press 2004) 231. This kind of reason is often distinguished from motivating reasons (ie those that agents actually act on, whether they *justify* their actions or not) and explanatory reasons (those describing why a decision, an action or any other phenomenon in general came about). For discussion, see John R Searle, *Rationality in Action* (François Recanati ed, MIT Press 2001) 97–134. For an interesting inquiry into the history of the distinction between normative and motivating reasons, see Jonathan Dancy, *Practical Reality* (OUP 2000) 20–6.

[34] In regarding the normative reasons of the kind, I call 'independent reasons' *facts*, following, among others, Dancy (n 33); Stephen Darwall, *Impartial Reason* (Cornell UP 1983); Joseph Raz, *Practical Reason and Norms* (OUP 1975); Scanlon, *What We Owe to Each Other* (n 33); Michael A Smith, *The Moral Problem* (Blackwell 1994).

[35] Rawls (n 29) 23: 'One is pictured as estimating on what percentage of the cases likely to arise a given rule may be relied upon to express the correct decision, that is, the decision that would be arrived at if one were to correctly [evaluate existing reasons]. If one estimates that by and large the rule will give the correct decision, or if one estimates that the likelihood of making a mistake by [evaluating the underlying reasons] directly on one's own is greater than the likelihood of making a mistake by following the rule, and if these considerations held of persons generally, then one would be justified in urging its adoption as a general rule.'

[36] Note that when Judge Easterbrook writes: 'Experts try for the best and most nuanced understanding; amateurs should use the approach that causes the least damage when it goes wrong. And judges are amateurs', he means that judges are not social scientists and must therefore rely on 'a relatively simple and mechanical approach to interpretation' or 'use an appropriately modest interpretive strategy [in light of the limitation in time and resources available to them]'. See Frank H Easterbrook, 'The Absence of Method in Statutory Interpretation' (2017) 84 U Chi L Rev 81, 83, 96–7. We can therefore assume that judges are 'experts' not in the sense of having social scientific expertise, but in the sense of having expertise in deploying ordinary tools of statutory interpretation, aside from the semantic canons but including 'relying on text', which Easterbrook maintains 'does the least harm'.

[37] For the view that reliance on the rules would be justified based on time- or labour-saving considerations, see section 6 below.

'that people are not able to [do the analysis] effortlessly and flawlessly' in each instance, Rawls explains.[38] Yet, given the relative fallibility of the underlying pattern, 'Each person is in principle always entitled to reconsider the correctness of [a summary rule] and to question whether or not it is proper to follow it in a particular case'.[39] If Sally were an expert, in the sense specified above, she should disregard the rule and determine whether Joe has complementary skills or not. This would ensure that she does not partner with an unskilled person in instances that fall outside the pattern underlying the rule. But since we are supposing that Sally cannot discern or weigh the independent reasons, that is, she cannot tell whether Joe indeed has complementary skills to her, she is entitled to rely on the rule. Even so, she cannot consider the rule as an *additional* reason to partner with Joe, meaning that it would be a mistake for her to think that there are *two* reasons to partner with Joe, namely Joe's likely complementary skills on the one hand and the said rule on the other. For the rule merely refers to the same underlying reason (the likelihood that Joe has complementary skills).

Thirdly, there may be countervailing reasons that the said rule leaves out. For, besides likely reasons to partner with Joe, which the rule captures, there are likely reasons against doing so that the rule does not capture. For instance, partnering with Joe may result in conflicts in executive decision making that prove fatal to the venture. What Sally should do is to weigh the reasons for and against partnering with Joe. And the existence of the tango rule, just as it does not bolster the reasons for partnering with Joe, does not defeat the reasons against it either. Thus, relying on the rule could result in her overlooking crucial countervailing reasons that bear on her decision. If Sally is not an expert in the sense above, she may not be able to discern these reasons either. Perhaps she has no choice but to follow the rule, risking both reliance on a fallible pattern and overlooking countervailing reasons. Regardless, she should not consider the rule as itself a reason that bears on the question. She should either close her eyes to the balance of reasons and rely on the rule or take the role of an expert and weigh the balance of reasons herself.

Of course, there are important differences between Sally and someone who struggles to discern statutory meaning, which differences I address in section 3. Nevertheless, Sally's case and the pitfalls that lie on her deliberative path helps us identify common mishaps in statutory interpretation.

Consider, for example, *Silvers v Sony Pictures*, which is a case widely known for its use of the *expressio unius* canon (literally, 'expression of one').[40] The court in *Silvers* treated that canon as though it were a reason in favour of the outcome propounded by the majority. It observed that the 1976 Copyright Act authorises a legal or beneficial owner of an exclusive right to sue for infringement of a copyright, further noting:

[38] Rawls (n 29) 23.
[39] ibid.
[40] See *Silvers v Sony Pictures Entm't, Inc* 402 F.3d 881 (9th Cir 2005) (*en banc*).

> The statute does not say expressly that *only* a legal or beneficial owner of an exclusive right is entitled to sue. But, under traditional principles of statutory interpretation, Congress' explicit listing of who *may* sue for copyright infringement should be understood as an *exclusion* of others from suing for infringement. The doctrine of *expressio unius est exclusio alterius* as applied to statutory interpretation creates a presumption that when a statute designates certain persons, things, or manners of operation, all omissions should be understood as exclusions.[41]

The court's majority here assumed the role of an expert—as it should—when it observed several factors as bearing on the question. Specifically, it relied first on the plausible presumption that standing to sue must be expressly granted by Congress, and secondly on textual evidence suggesting that 'Congress' grant of the right to sue was carefully circumscribed' even for those expressly included in the statute.[42] Yet, it considered these reasons *along* with the canon, as though the existence of the canon was itself a reason in favour of its adopted interpretation or that it somehow added weight to such reasons. The majority also erred in thinking that the existence of the canon somehow nullifies or renders irrelevant the countervailing reasons in the act's legislative history.[43] Instead of balancing competing independent reasons—say, m reasons in favour and n against their adopted interpretation—the majority relied on some of the independent reasons, namely the m reasons in favour, regarded the canon as yet an m + 1th reason and ignored countervailing independent reasons (assuming that the existence of the canon somehow renders the sum of reasons in favour of its adopted interpretation conclusive). I have not yet argued that canons are summary rules (that will be done in section 3). But if the *expressio unius* canon is a summary rule, then the court's analysis is mistaken insofar as it treats that canon as a reason that could support its adopted interpretation or defeat evidence against it.

The scholarship on the semantic canons frames questions of statutory interpretation around the potential application of canons and counter-canons in every case. Sometimes, statutory litigation and judicial decisions are also framed around competing rules.[44] It is therefore helpful to examine cases of conflicting summary rules. To that end, let us consider Sally's case, but this time with duelling summary rules. As discussed, there are likely reasons for and against her partnering with Joe. Consider now four possible worlds that are in all respects identical, except with regard to what summary rules exist for Sally's consideration.

In possible world one, she does not know any summary rules and must decide without them.

In possible world two, before she makes her decision, she is taught the maxim 'two sailors sink a ship'. It is clear that she would have no more reason against the partnership in world two than she has in world one. If, prior to learning the rule, she had m reasons against the partnership, she would not now, in light of the

[41] ibid (internal quotation and citation omitted, original emphasis).
[42] ibid 885.
[43] I discuss legislative history in s 3B.
[44] For data on this point, see Anita S Krishnakumar, 'Dueling Canons' (2016) 65 Duke LJ 909.

rule, have $m + 1$ reasons against it. If the reasons Sally has on each side amount to 50–50, being reminded of the 'two sailors' rule will not shift the balance to 49–51 against partnering.[45]

Now suppose that in possible world three, Sally is instead taught the maxim 'it takes two to tango' before making her decision. It is equally clear that she does not thereby gain additional reasons in favour of the partnership in this world. This summary rule does not *weigh in* as a new reason in her deliberation either.

Finally, suppose that in possible world four, Sally is taught both rules before deciding. Once again, the reasons Sally has for and against partnering with Joe remain the same. Thus, the existence of summary rules does not play any role in determining the outcome to Sally's deliberation.

Now we can finally cast our example in terms of a duelling canon's scenario that permeates the scholarship's imagination, not least due to Llewellyn's famous characterisation. According to that scenario, the outcome of Sally's deliberation would be determined by weighing the rules against one another. As discussed above, Llewellyn only rejected this approach because he saw no way to adequately weigh the rules. Since then, much scholarly wrangling has focused on the right criteria to weigh duelling rules. But our example clearly shows that Sally's dilemma is not one of duelling rules. She should not, in other words, think: 'What am I to do? On the one hand, there is a rule in favour of partnering with Joe, but on the other hand, there is another rule against it.' Instead, she should think something like: 'What am I to do? On the one hand, having a trustworthy partner is an asset, but on the other hand, we may get into disagreements. On the one hand, Joe is very good in getting us the right deals, but on the other hand, he may mess up the accounting.'

And so forth. The correct outcome to her deliberation is entirely dependent on the relative weight of such reasons, which means that whether she should partner with Joe or not can and must be determined independently of any summary rules whatsoever.

If my argument (in section 3) that canons are summary rules proves right, then Sally's example shows precisely in what way the scholarship follows Llewellyn's seminal work on the wrong foot: questions of statutory interpretation are not questions about competing summary rules; they are questions about statutory meaning which can and should be resolved independently from such rules.

Summary rules are not normative because they are *post hoc* generalisations. The observations we made in the thought experiments are not contingent on the presupposed conditions that govern them. Rather, as Rawls points out, they occur systematically due to a defining feature of summary rules: summary rules are *post hoc* generalisations about features of decisional instances, where the instances are logically prior to the rule.[46] They are conceptual representations of patterns that we can only discern in hindsight. The reason that summary rules do not tip the

[45] Similarly, if her reasons for deciding to partner, for example, outweighs the other reasons, they will not be weakened if the 'two sailors' rule is invoked. This is to be expected, given the 'two sailors' rule is a summary rule.

[46] Rawls (n 29) 22.

balance in particular instances is that they merely track the reasoning behind an aggregate number of other (past or hypothetical) individual decisions. As such, regardless of how accurately they track the reasons in those other cases, they cannot weigh in on whether any particular decision should accord with them or not. Thinking that the instance should comply with the rule based on the accordance of the rule with similar past or hypothetical cases amounts to confusing the normative causes with effects: how a decision is to be made in any particular instance will affect the content and accuracy of the rule, but the rule's content has no normative implications for how an instance is to be decided.

3. The Semantic Canons Do Not Weigh in on Statutory Interpretation

Thus far, I have shown what summary rules are and why they do not weigh in on practical deliberation. I have also suggested that, in the same way, the semantic canons do not weigh in on questions of statutory interpretation. But I have not yet addressed whether the semantic canons are summary rules. Nor have I addressed in what sense the semantic canons in the context of statutory interpretation are akin to summary rules in practical deliberation. This is because my discussion thus far has revolved around Sally, whose situation may be reasonably regarded as sufficiently different from someone trying to decipher statutory meaning.

I shall now address these matters, first arguing that the semantic canons are indeed summary rules (section 3A), then discussing the similarity of their role in questions of statutory interpretation to the role summary rules play in Sally's deliberations (section 3B).

A. Semantic Canons are Post Hoc Generalisations

Establishing that the semantic canons are summary rules requires an argument that shows that interpretative decisions about statutory language are logically prior to the semantic canons. In other words, I must show that it is not that instances of statutory interpretations get their correctness by way of according with a semantic canon, but that a semantic canon gets *its* validity from the fact that it expresses a pattern of statutory interpretations that are independently correct in numerous prior instances. I can do this by showing that the semantic canons are *post hoc* generalisations that play no role in the correctness of various instances of statutory interpretation. This will mean that the semantic canons can never be the *reason* for the correctness of an interpretation, or equally, that no interpretation owes its correctness to its accordance with the canon *qua* a rule.

Such an argument was offered by Lewis Carroll back in 1895, although it did not target the semantic canons directly.[47] Carroll, who was equally renowned as a logician as he was as a writer of children's fiction, made such an argument

[47] Lewis Carroll, 'What the Tortoise said to Achilles' (1895) 4 Mind 278.

regarding the rules of formal logic. More precisely, by way of formulating a paradox, Carroll pointed out that the rules of formal logic play no role in the validity of logical inferences.

The paradox is introduced in a playful dialogue between Achilles and a tortoise regarding the following argument:

(A) Things that are equal to the same are equal to each other.
(B) The two sides of the triangle are equal to the same.

Therefore,

(Z) The two sides of the triangle are equal to each other.

The tortoise accepts the premises (A) and (B), but denies the conclusion (Z), asking Achilles to show how he would be logically forced to accept (Z). Achilles assumes that by getting the tortoise to accept the rule of *modus ponens*, namely that 'if (A) and (B) are true, then (Z) must be true', he can force the tortoise to accept (Z). The tortoise grants the rule, and asks Achilles to write it down as an additional premise, calling it (C):

(C) If A and B are true, Z must be true.

But the tortoise still asks to be shown what forces him to accept (Z), thereby making Achilles provide *ad infinitum* such additional premises:

(D) If (A) and (B) and (C) are true, Z must be true.
(E) If (A) and (B) and (C) and (D) are true, Z must be true.
…

This infinite regress arises due to Achilles's *reliance* on the rule of *modus ponens*. Let me unpack this.

Much that has been written on this paradox takes it that the regress can be prevented if, at any stage, the tortoise *applies* the rule of *modus ponens* to the premises he accepts.[48] On these accounts, Achilles' mistake was to take *modus ponens* as a missing *premise* of the argument rather than a *rule*.[49]

However, even if we suppose that *modus ponens* is a rule rather than a premise, the infinite regress persists, because the tortoise has no reason to apply *modus ponens* unless, of course, there was another rule that mandated *its* application, and so forth with each rule being in need of yet another.[50] The fact that this path, too, requires infinitely more rules shows that the validity of (Z) cannot rely on the *application* of any such rule either.

The only way to avoid the regress is to deny the first step that leads to it, namely, Achilles' presupposition that the rule plays any role whatsoever in the

[48] Nico Kolodny and John Brunero, 'Instrumental Rationality', *Stanford Encyclopedia of Philosophy* (Winter edn, 2016) <https://plato.stanford.edu/archives/win2016/entries/rationality-instrumental>.

[49] See Peter Railton, 'On the Hypothetical and the Non-hypothetical in Reasoning about Belief and Action' in Garrett Cullity and Berys Nigel Gaut (eds), *Ethics and Practical Reason* (OUP 1997) 76–9.

[50] As mentioned in the text accompanying n 28 above, Manning and Stephenson (n 4) 214 foresee such a problem with respect to the canons of interpretation.

validity of (Z).[51] To avoid the regress, we must note that the inference of (Z) is valid as it stands and not in virtue of any additional rules.[52] This validity is owed solely to the entailment of (Z) by the premises (A) and (B), and not to any external premises or rules.

Carroll makes this point by showing that a failure to see (Z)'s entailment in (A) and (B), ie (Z)'s denial while simultaneously accepting (A) and (B), would not be corrigible by the invoking of *modus ponens*. This is because one can only take *modus ponens* to be a correct rule insofar as one takes inferences like (Z) as valid. From a position that (Z) is not entailed by (A) and (B), the instance of (Z) counts as a counterexample to the rule of *modus ponens* (and grounds for the rejection of it).[53] Thus, *modus ponens* could not help someone to arrive at (Z) who does not already take (Z) as entailed in (A) and (B). This is the position that the tortoise takes, namely that if the entailment of (Z) in (A) and (B) were to be challenged, the rule (C) could not help.

Note that the tortoise accepts the rule of *modus ponens* by granting (C) but denies that it can necessarily be said of this instance and thereby compel the validity of (Z). Nor can the validity of (Z) be explained in terms of the argument's accordance with the rule. In short, to suppose that it is thanks to *modus ponens* that (Z) is entailed by (A) and (B) is to already put the cart before the horse.

The argument offered by Lewis Carroll is much stronger than what would serve our purposes here, which would be served by merely showing that *modus ponens* is a summary rule. This amounts to showing that *a tortoise who knows the rule of* modus ponens *has no more reason to accept (Z) than another tortoise who does not know such a rule*. A tortoise in a world without the rules of formal logic would be in the same position.[54] This can be easily gathered from the resolution to the paradox offered above, namely that (Z) can be validly inferred without ever referring to *modus ponens* or any such rule whatever. And the same goes for all rules of formal logic, because they only name a pattern of syntactical form that arguments whose conclusions are entailed by their premises take.

What do the rules of logic such as *modus ponens* have in common with the semantic canons? They are all names for syntactical patterns that occur in numerous actual or hypothetical instances and are therefore all logically *post hoc* to the particular instances of their applications. Specifically, the rules of logic are formalisations of syntactical patterns that occur in an infinite number of semantic inferences that are independently valid. Their only difference from the semantic canons is that they express a formal syntactical pattern that *always* coincides with the independent characteristic of interest, namely the entailment of the conclusion by its premises. In other words, the rules of logic are *perfect* generalisations,

[51] Searle (n 33) 19.

[52] ibid.

[53] See ibid: 'It would be more accurate to say that the rule of modus ponens gets its validity from the fact that it expresses a pattern of an infinite number of inferences that are independently valid.'

[54] I do not mean to argue that there may be a world where the generalisations expressed in the rules of formal logic fail, but only that neither the fact that such generalisations obtain nor the fact that they are expressed in the form of formal rules tips the balance of reasons for the tortoise.

in the sense that they are without exceptions. But there are plenty of other syntactical patterns that can be, and indeed are, *imperfectly* generalised. Logical fallacies are of this kind: they express a syntactical form of argument that only sometimes coincides with their independent characteristic, namely the entailment of the conclusion by the premises.

The semantic canons are members of this latter category of generalisations, meaning that they too only sometimes coincide with their respective semantic characteristics of interest.[55] Both logical fallacies and semantic canons are *the same type* of formalisation as the rules of logic, their differences notwithstanding, and are therefore equally subject to the Lewis Carroll paradox. To see this, consider the following invalid inference:

> If Nelson is a human, then Nelson is mortal.
> Nelson is not human.
> Therefore, Nelson is not mortal.

The argument here is said to be an instance of the fallacy of denying the antecedent. That logical fallacy is the form of argument in which the argument's second premise denies the antecedent of the hypothetical in its first premise and concludes the negation of its consequent:

> The fallacy of denying the antecedent: $((p \rightarrow q) \ \& \ \sim p) \rightarrow \sim q$.

The argument in any instance of this fallacy is said to be invalid, meaning that it is possible for its conclusion to be false even if all its premises are true (after all, Nelson could be a cat).[56] But this possibility does not arise from the argument's accordance with the fallacy, which formalises the pattern of false conclusions in an infinite number of other arguments with a similar syntactical form. Rather, it arises from the fact that the premises of *this* argument do not entail its conclusion. As such, it is the pattern in question that arises from an infinite number of independently invalid instances such as the one at hand. Therefore, it would be a mistake—in fact, the same mistake as the one that Achilles made—to suppose that the conclusion of the argument may be false *because* it is an instance of the fallacy. The conclusions of this and many other such arguments may be false *as they stand*, and not due to any pattern that may appear in their syntactical form.

I discuss the fallacy of denying the antecedent because of its affinity with the *expressio unius* canon (also known as *inclusio unius*), which we encountered in the discussion of *Silvers* above. Recall that this canon holds that if a statutory provision explicitly names particular elements as within its scope, other elements that are not mentioned are implicitly excluded from the scope of its application. This canon is a textbook example of the fallacy of denying the antecedent, if not

[55] This makes the results of relying on them as unreliable as inferences made based on logical fallacies. But that is not at all the point I wish to make.

[56] Biographical note: indeed, Nelson is a cat. Named after Admiral Nelson of the British Royal Navy, he also boasts a fleet of six Instagram followers.

a Latin name for it. It states: if an element is expressed, then it is included; an element is not expressed; therefore, it is not included.

The *expressio unius* canon: $((p \rightarrow q) \ \& \ \sim p) \rightarrow \sim q$.

It is important to note that the point of the current discussion is *not* to criticise the canon for being fallacious. The point here is that the canon, *qua* a rule, bears no normative force in settling questions of statutory interpretation. What follows from this point is that the correctness of a resolution to such a question depends purely on factors that are independent from semantic canons. With this in mind, let us revisit *Silvers*.

In that case, competing reasons divided the court 4:5.[57] On the one hand, two reasons seemed to support the majority's finding: first, the presumption that standing to sue must be expressly granted by Congress; and second, textual evidence suggesting that even rights that were expressly granted in the statute were carefully circumscribed. On the other hand, the statute's legislative history seemed to weigh against the majority's finding.[58]

What the court should have done was to weigh these reasons against one another. The *expressio unius* canon is simply a label for a syntactic form that expresses the exclusion of elements not explicitly included in countless (though not 100% of) cases. However, in none of those cases is it the syntactic form in which the idea is expressed that excludes those elements.[59] Yet, the majority dismissed legislative history by appeal to the canon. Meanwhile, in his dissent, Judge Bea retorted that there's a duelling canon that could be used, that the canons are not binding and that they should only be considered if legislative intent could not be determined.[60]

B. Statutory Interpretation as Deliberation

We saw that summary rules that apply to Sally's case do not tip the balance of reasons in her deliberations. What must be true for the semantic canons to similarly play no role in questions of statutory interpretation? I want to suggest that, insofar as interpretive deliberation can be analogised to practical deliberations, semantic canons—and syntactical rules more generally—are the analogues to summary rules.

In practical deliberations, one struggles to determine what she ought to do. Certain factors will bear on the outcome, so that, all things considered, their

[57] Judge Berzon and Judge Bea each wrote a dissent, with Judge Reinhardt joining in Berzon's dissent and Kleinfeld joining Bea's dissent; *Silvers* (n 40).

[58] ibid 899–900 (internal quotation and citation omitted, original emphasis).

[59] To emphasise, this is not because we are dealing with a fallacy, but because that form says nothing about the (in)validity of the inferences. Thus, even if, rather than the fallacy of denying the antecedent, we gave a rule of logic, say *modus ponens*, a place among the semantic canons, it too would never be what makes the interpretive inferences it labels valid. Recall that a tortoise who knows the rule of *modus ponens* has no more reason to accept (Z) than another tortoise who does not know such a rule. In the same way, a judge who knew the canon of *modus ponens* would have had no more reason in support of the candidate interpretation that rule predicts than one who did not know such a canon.

[60] *Silvers* (n 40) 899–900 (internal quotation and citation omitted, original emphasis).

convergence determines the correct outcome of practical deliberation. Call these factors normative determinants. Normative reasons in the sense specified above, that is, considerations in favour of or against certain actions, are among the normative determinants of deliberations. Philosophers often disagree about *what* factors these normative determinants are. Rules of morality, rationality and prudence are some plausible but by no means uncontroversial examples. That is to say, philosophers disagree about whether the correct output of practical deliberation is sensitive to the content of such rules, though most take such rules as at least initially plausible potential candidates. Said disagreements, however, need not worry us. For our purposes, we need not specify what normative determinants are. All we need to do is to call whatever factors that do in fact determine correct outcomes of practical deliberations 'normative determinants'.

What the discussion of summary rules and the examples of Sally in this article show is that summary rules are *not* normative determinants. Instead, they are labels for robust patterns of the existence of normative determinants within the context of practical deliberations.[61] Since summary rules portend the existence of normative determinants, we can call them 'determinant predictors'. As we saw, the determinants that summary rules portend are not necessarily conclusive but must be weighed against other determinants that may exist. Moreover, because the determinants that summary rules portend are not always present, we can call them 'imperfect determinant predictors'.

In interpretative deliberations, one struggles to understand the meaning of some text. Certain factors will bear on that outcome as well, so that the correct meaning of a text will be the convergence of all such factors. Call these factors *semantic* determinants. Philosophers also often disagree about what semantic determinants are. Authorial intentions (speaker-meaning), common usage, pragmatic considerations, circumstantial context, textual context (what else was said) and so on are some plausible but controversial factors. That is to say, philosophers disagree about whether the correct output of interpretive deliberation is sensitive to such factors, though most regard them as at least initially plausible potential candidates.

My argument above suggests that in interpretive deliberation, the category of rules to which rules of logic, logical fallacies and the canons belong is the analogues to summary rules in deliberative interpretation. This means that insofar as the two deliberative processes are analogues, syntactic rules must be considered determinant predictors rather than semantic determinants. Since canons and fallacies make fallible predictions, they too are 'imperfect determinant predictors' (as opposed to rules of logic, which make necessary predictions and are therefore 'perfect determinant predictors'). It is this analogy between the two activities that allows me to conclude that the semantic canons do not tip the

[61] Thus, as the examples suggested, when I say that the 'tango' rule applies to Sally's case, what I mean is that in the context of her deliberation, there is a robust pattern in the existence of certain normative determinants.

balance of semantic determinants in questions of statutory interpretation just as summary rules do not tip the balance of normative determinants in Sally's practical deliberation.

My analysis says nothing about what semantic determinants are, given I need not take a stance on this matter to make my point. All I can say is that canons are not semantic determinants. Presumptions about consistency and non-redundancy, legislative intent, legislative context or even historical context are all potential candidates that initially appear plausible but are far from uncontroversial. But, given that my critique operates in a narrow scope, I must stay out of these debates. The narrow scope of my critique leaves all such factors unscathed; that is, strictly in the view of my argument here, all such factors are presumptive candidates of being semantic determinants. These presumptive candidates may be defeated on other grounds besides mine, but they cannot be defeated by the canons.

Recall once again the issue in *Silvers*: the purported reasons against the majority's finding were said by the dissent to lie in legislative history. Whether legislative history amounts to such reasons at all is itself a matter of heated debate. Suppose that an analogue argument to the one in this article suggests that legislative history is also no semantic determinant (at least, not in the context of statutory interpretation).[62] Such an argument could defeat legislative history in general and would have made appeals to it in *Silvers* moot. But an appeal to the *canons* could not have done that, at least not in light of my argument here. Now recall that this was exactly what the majority opinion tried to do. It is this move of the majority that entitles me to criticise them for disregarding legislative history, while remaining agnostic about whether outcomes of statutory interpretation are sensitive to legislative history at all. In other words, I can readily grant that the majority could have reasonably disregarded legislative history in light of the postulated analogue to my argument here, while criticising them for dismissing legislative history by appeals to the canons instead. Once again, the issue I raise is with the court's rationale rather than its outcome.

4. What's Sauce for the Goose Is Sauce for the Gander

What goes for summary rules goes for the semantic canons, though not because of the maxim invoked in the title of this section. This interpretive maxim has no bearing on this question. What goes for summary rules goes for the semantic canons because the semantic canons are a species of the summary rules genus and therefore imperfect determinant indicators by nature. As determinant indicators, the semantic canons do not tip the balance of reasons in questions of statutory interpretation.

Judges, practitioners and commentators, however, mistakenly regard the canons as normative determinants, that is, rules that provide reasons that can tip the

[62] For a recent defence of using legislative history in statutory interpretation, however, see Robert A Katzmann, *Judging Statutes* (OUP 2014) 39.

balance of reasons. This is what they mean when they say that even if they are not dispositive, the canons support their choice of interpretation. As evident in my discussion of *Silvers* above, they weigh the canons not just against one another (in a duelling scenario), but also against other reasons (such as legislative history). They take the canons as one of a number of arguments in favour of a candidate interpretation. They think that if, besides the canons, there are *m* reasons in favour and *n* against their choice of interpretation, then with a 'supporting' canon, there will be *m* + 1 in favour and *n* against, and if there is a duelling canon too, then there will be *m* + 1 for and *n* + 1 against.

My analysis shows that this is mistaken and that these rules should not be treated as reasons. It also clarifies that lasting debates about the relative strength of the canons, how they measure against other interpretive tools or how they are to be weighed against one another are all misguided. They are misguided because they are based on the misconception that the canons have some, albeit hard to determine, weight or force. Once we liberate ourselves of that misconception, the issues underlying these debates dissipate.

To see just how vexatious this misconception can be, consider *People v Smith*, in which the statute at issue enumerated dagger, dirk and stiletto as examples of dangerous weapons it barred.[63] The court invoked *ejusdem generis* and limited the scope of the statute to 'stabbing weapons'.[64] Commenting on *Smith*, Manning and Stephenson note that relying on *ejusdem generis* alone does not make clear which of the shared characteristics of the terms in the list defines the type to which the generic term following the list ought to belong.[65] They remark that the *Smith* court could 'just as easily have defined the *ejusdem generis* category as "weapons designed for use in hand-to-hand combat"'.[66] Given what was said, we can unpack the rationale of the court and Manning and Stephenson's worry.

First, the court's decision did not solely rely on the canon. Instead, the court offered independent analysis to show that, given the meaning of the relevant portions of the statute, the M-1 rifle that the defendant was found with could not count as a dangerous weapon. For example, the court pointed out that other parts of the statute made carrying weapons like M-1 rifles illegal only when a culpable *mens rea* accompanied the act.[67] In addition, the court noted that interpreting that portion of the statute as covering M-1 rifles would have made many residents who went hunting potentially liable to a felony and, as a result, inferred that covering M-1 rifles could not have been the legislators' intent.[68] The court's determination that M-1 rifles did not fall within the scope of the generic term in this case was based on such independent reasons, though a canon was nevertheless invoked.[69] Thus, though commentators often report the court's decision as being based on the canon, this was not strictly so.

[63] *People v Smith* 393 Mich 432, 435 (1975).
[64] ibid 436.
[65] Manning and Stephenson (n 4) 236.
[66] ibid.
[67] *Smith* (n 63) 437–8.
[68] ibid 437.
[69] ibid 436, 437.

Nevertheless, Manning and Stephenson are right, because the court's decision did not stop there. Instead, it further specified a dangerous weapon as 'a *stabbing* weapon'.[70] Manning and Stephenson rightly worry that this was an unjustifiable overspecification of the law. But the problem is not so much that the specification was too narrow to be justifiable by the canon alone, as Manning and Stephenson say.[71] Rather, the problem lies in thinking that the correct specification must be discerned through the canon's lens. Let me unpack this latter point.

Note that all that was needed to decide the question raised in *Smith*—namely whether an M-1 rifle was a dangerous weapon within the meaning of the relevant section of the statute in question—was the independent analysis that the court had already offered. Why, then, did the court think that a further specification of what 'dangerous weapons' are is necessary at all? I think because the court was thinking more about the canon than the specifics of the case. The case did not require the specification of the items enumerated in the statute in under any more specific category than 'dangerous weapons'. What required the further specification of that category was the *ejusdem generis* canon. It is the *canon* that motivates discussion of unifying features. Thus, the court went searching for a characteristic, common among the elements of the list, to serve as the defining feature of the class to which they and the generic term belong, likely only because it looked at the case through the lens of *ejusdem generis*.

Consider now Manning and Stephenson's retort that relying on *ejusdem generis* alone does not make clear which of the shared characteristic of the terms in the list defines the type to which the generic term following the list ought to belong.[72] They worry that the court's specification was too narrow, while completely overlooking the fact that it was not at all necessary. They then go on looking for what categories satisfy the canon, noting that the *Smith* court could 'just as easily have defined the *ejusdem generis* category as "weapons designed for use in hand-to-hand combat"'.[73] True, once feet are set on the path of searching for what satisfies the canon, electing an arbitrary characteristic as the elements' singularly unifying feature becomes unavoidable. But the problem lies in the misconception that questions of statutory interpretation are questions along that path.

5. *What if the Legislature Legislates with the Canons in Mind?*

One of the background questions for statutory interpretation is whether the legislature legislates with the semantic canons in mind. Most empirical studies answer

[70] ibid 436 (emphasis added).
[71] Manning and Stephenson (n 4) 236.
[72] ibid 236.
[73] ibid.

this question in the negative.[74] But would anything change had that answer been empirically positive? The facts notwithstanding, it is at least theoretically possible that the legislature, at least partly, considers the canons when legislating. Would such a fact change the role of the canons? The answer to this question is more complicated than may seem.

First, it is important to understand that in the case that the legislature considered the canons when legislating, all that the canons could have revealed would have been *legislative intent*. Yet, as *Silvers* shows, actual proponents of the canons are often *textualists* who state their aim not as reviving congressional intent, but as making the judicial process itself more efficient and objective *without* having to reconstruct the legislative intent.[75] Thus, the actual proponents of the canons are barred from such a line of defence for relying on them. It is important, therefore, to note from the outset that the potential substantive role of the canons in our counterfactual thought experiment here would be the subject of an entirely theoretical musing rather than a viable argument in favour of relying on the canons. Equally, none of the proponents of consulting legislative intent in statutory interpretation regard the canon as a reliable way of doing so. Nor would the canons be the best source of conveying legislative intent when there is legislative history.

The implication here is that even if the legislature legislated with the canons in mind, there may well be no uptake on the side of the interpreters. For communication, we need both a speaker and a listener, an author and a reader, a legislator and an interpreter. Supposing that the legislature legislates with the canons in mind by itself only ensures one half of that process. To suppose that legislative intent can be communicated through the canons, we must also suppose that judges interpret legislation with the canons in mind.

That too, however, conflicts with the empirical evidence. As Abbe Gluck notes, the resistance of judges to a formalist application of semantic canons and other tools of statutory interpretation 'seems to stem from a constitutional law-level intuition that choice of interpretive method is inherent in each individual judge's power to judge'.[76] Gluck not only finds that judges reserve such a right, but also that they regard having such a discretion as a matter of constitutional

[74] See, eg Abbe R Gluck and Lisa Schultz Bressman, 'Statutory Interpretation from the Inside—An Empirical Study of Congressional Drafting, Delegation, and the Canons: Part I' (2013) 65 Stan L Rev 901 (showing that Congress does not always know about or follow the assumptions judges make about legislative drafting and interpretation); Lisa Schultz Bressman and Abbe R Gluck, 'Statutory Interpretation from the Inside—an Empirical Study of Congressional Drafting, Delegation, and the Canons: Part II' (2014) 66 Stan L Rev 725. For discussion on *noscitur a sociis* and *ejusdem generis*, see Gluck and Bressman 933. For similar findings in agency rule-drafting practices, see Christopher J Walker, 'Inside Agency Statutory Interpretation' (2015) 67 Stan L Rev 999, 1026–7 (showing that although a majority of drafters use concepts such as *noscitur a sociis* and *ejusdem generis*, such concepts are not always used in practice).

[75] On this point, see also Frank H Easterbrook, 'Text, History, and Structure in Statutory Interpretation' (1994) 17 Harvard Journal of Law & Public Policy 61, 63. However, see Anita S Krishnakumar, 'Backdoor Purposivism' (2020) 69 Duke LJ 1275.

[76] Abbe R Gluck, 'Justice Scalia's Unfinished Business in Statutory Interpretation: Where Textualism's Formalism Fell Short' (2017) 92 Notre Dame L Rev 2053 (adding that 'An important ancillary consequence of the judicial unwillingness to follow through with formalism is that Congress views the courts' interpretive rules as arbitrary and results oriented, and so not worth learning or coordinating with'); see also Gluck, 'Congress, Statutory Interpretation, and the Failure of Formalism' (n 8) 184–5.

law. This finding is consistent with what judges report in surveys[77] and in judicial scholarship.[78] So long as judges regard themselves as having discretion to question whether it is proper to follow a canon in a particular case, it is hard to see what the theoretical possibility of legislatures legislating with the canons in mind would change.

Nevertheless, it is also within the realm of theoretical possibility that judges give up said discretion (the legality of the move aside and the disastrous inaccuracies that would result notwithstanding). Thus, we can revise our hypothetical to suppose that not only the legislature considers the canons when legislating, but that the courts also interpret legislation with them in mind. In such a world, though it may be distant from ours, the canons may communicate, albeit haphazardly, legislative intent.

6. Can the Canons be Used as Time- or Labor-Saving Devices?

Although we may accept that summary rules lack normative force as a theoretical matter, we must admit that in most real-world scenarios, decision making involves transactional costs. Once such costs are taken into account, summary rules seem to only lack normative force for decision makers capable of discerning and weighing all existing reasons in addition to having enough time so that time is not a constraint on deliberation. Surely that is not the case in a typical instance of statutory interpretation. Do such considerations justify judges in relying on the semantic canons as time- or labour-saving devices?

The idea of using summary rules as time-saving devices goes at least as far back as JS Mill,[79] but its most systematic presentation has been made by Joseph Raz. On Raz's view, rules are justified as time-saving devices, labour-saving devices and devices to reduce the risk of error in deciding what ought to be done where we cannot trust our judgment in estimating them.[80] Such rules are contemplated in tranquility to state what is to be done on the balance of foreseeable reasons in situations to which they apply. When a situation to which such a rule applies actually occurs, one can rely on the rule, 'thus saving much time and labour and reducing the risks of a mistaken calculation which is involved in examining afresh every situation on its merits'.[81] However, this kind of usage is only conditionally sanctioned.

Raz for example, only sanctions reliance on such rules if the rule (i) has been contemplated in light of the balance of the totality of foreseeable reasons and

[77] Abbe R Gluck and Richard A Posner, 'Statutory Interpretation on the Bench: A Survey of 35 Judges on the Federal Courts of Appeals' (2018) 131 Harv L Rev 1298 (showing that most judges place a high value on retaining some flexibility in statutory interpretation).

[78] See eg Easterbrook, 'The Absence of Method in Statutory Interpretation' (n 36) (showing that the courts do not always follow the same rules).

[79] John Stuart Mill, *A System of Logic* (1843).

[80] Raz (n 34) 59.

[81] ibid.

(ii) reliance on it makes one better off. Starting with the second condition, Raz's demand is not that the rule always correctly represents the balance of foreseeable reasons, but that it does so more often than direct judgment would. Thus, the second condition imposes substantive limits on the relatively higher accuracy of the rule over direct judgment. The first condition imposes a procedural limitation by ensuring that the rule is of a type that can, at least theoretically, hope to reasonably satisfy the second condition. Thus, the first condition, though being the mechanism for guaranteeing higher substantive accuracy of the rule, constrains its justification by something beyond its mere accuracy. This constraint would disqualify arbitrary rules that systematically leave out important reasons or take into account imaginary ones, even if, by some miracle, they achieved higher accuracy. Put differently, Raz's framework is concerned not merely with whether the rules produce more or less positive outcomes, but also with their potential for rational justifiability.

For example, we may design a set of rules for manufacturing a product in mass quantities without having to design each individual unit of that product separately. In such a case, all the relevant considerations in production would be reflected in said rules and reliance on them comes at no deliberative loss. Whether each unit is designed separately or following rules for mass production, the outcome remains the same.[82] Contrast this to the duelling rules of thumb that applied to Sally's deliberation. In that situation, each of the two rules leaves significant considerations out of sight. Neither of them can be sanctioned in the same way because (i) relying on them will likely lead Sally to make more mistakes than direct judgment would and (ii) the rules are not even theoretically apt to lead Sally to right decisions because they only contemplate a subset of potential reasons and therefore cannot even hope to represent the balance of all foreseeable reasons.

Let me unpack this latter point. First, I made a distinction between foreseeable and potential reasons. In the case of mass-producing units, the rules contemplate foreseeable reasons, that is, reasons that will actually exist in ordinary circumstances. Not so in Sally's case. As discussed, we can imagine a scenario in which Joe does not have the expertise we supposed he does. Even in such an ordinary circumstance, the 'tango' rule represents reasons that are not at all present. Thus, whereas the production rules signal reasons that will exist under ordinary circumstances (ie foreseeable reasons), each of the two rules that apply to Sally represents reasons that only *may* exist under ordinary circumstances.

Secondly, I distinguished rules that contemplate a subset of reasons from those that contemplate all reasons. Sally's rules, even if they represented foreseeable

[82] For simplicity, I am assuming that production rules are accurate in predicting foreseeable reasons. But this assumption need not be made for the argument to hold. Rules like production rules are also only as accurate as the statistical regularity of their underlying patterns. The point I am trying to make is that canons with even perfect statistical regularity in their underlying patterns still only tell one side of the story. Put differently, whereas underlying production rules is a pattern in the balance of all reasons, underlying the canons are patterns in the existence of raw reasons yet to be weighed against others.

rather than potential reasons, represent only a subset thereof rather than their balance. Production rules, on the other hand, *weigh all existing reasons and represent their balance*. Sally's rules each represent some reasons and leave it to Sally to weigh them up. Sally's reliance on either of the rules would wholly lack rational justification (of the kind production rules enjoy) in addition to likely leading to extremely poor decisions.

As the analysis in previous sections shows, the canons are akin to Sally's rules rather than production rules. They represent a subset of reasons that may exist in any given case of the same type. That the canons represent potential rather than foreseeable reasons underlies the need for judges to decide whether a canon or its counter-canon applies at all. That they represent a subset of reasons rather than the balance of their totality underlies the need for judges to choose between a canon or its counter-canon. In some ways, we can see where the thrust of Llewellyn's critique comes from: the canons are not rules that judges can merely follow; they are rules that may or may not apply, and when they do they point in different directions. Llewellyn worried that there is no way to weigh the rules against one another, and my discussion shows exactly why: what needs to be weighed are not the rules, but the reasons underlying each rule and their relative weights. However, my analysis further shows that that activity can be done entirely without the canons. Any thought or talk of the canons will result in one or both of the following mistakes: canons get mistaken as rules that are supposed to replace direct judgment (like production rules) or as amounting to tipping reasons themselves.

Going back to the question we started with, namely whether the canons can be used as time-saving devices: first, since the underlying potential reasons the canons represent must themselves be weighed, no actual time will be saved. In other words, using the canons as time-saving devices would necessitate reliance on them in lieu of directly assessing the actual balance of reasons. As Raz explains, factors like time- and labour-saving are not reasons that directly weigh in the deliberation about what is to be done, but are reasons to disregard the balance of reasons and to act on a predetermined rule instead.[83] In this sense, these factors are 'exclusionary reasons',[84] understood as reasons to disregard the balance of reasons.[85] Thus, treating rules as time-saving devices requires regarding them as exclusionary reasons: not as reasons that tip the balance of reasons, but as rules that supplant the balance of reasons. And if this is what is meant by 'using the canons as time-saving devices', then the answer is no, because, unlike production rules, the canons are not even theoretically apt as candidates for replacing the balance of reasons. If anything else is meant by 'using the canons as time-saving devices', then the answer is this: so long as directly weighing the reasons remains necessary, no time will be saved.

[83] Raz (n 34) 59–62.
[84] ibid 59–62.
[85] ibid 39: 'A second-order reason is any reason to act for a reason or to refrain from acting for a reason. An exclusionary reason is a second-order reason to refrain from acting for some reason.'

Judges regard themselves as having the task of (and the competence in) deciding whether a canon applies to a case or choosing between several canons that purportedly do apply. This is the purported task that has famously (and incorrectly) been explained as making a choice between duelling canons (*à la* Lewellyn) or between a canon and its exceptions (*à la* Lewellyn's critics, like Scalia). If the courts have this task, then there is no gain in efficiency because deciding whether canons apply at all and settling which among several potential candidates do both require working out the balance of underlying reasons. Think of Sally's dilemma: the only way she could resolve her 'duelling rules' scenario was by looking at the underlying reasons and weighing those out. But if she is torn between two rules and decides which to follow by considering all the underlying facts (like Joe's actual skill set), then she really does not save any time or epistemic labour. So, too, having to evaluate the underlying reasons means that the courts could not actually save any time or labour. Remember that saving time or labour is only possible when the balance of reasons is *disregarded* in favour of relying on the rule (like in the production case), but the balance of reasons could not be disregarded when the agent is faced with the task of choosing from among multiple rules that only potentially apply.

I want to emphasise that using the canons as time- or labour-saving devices is not an actual underlying issue in scholarship about the canons. Nor do courts treat the canons as such devices. Actual proponents of the canons speak of them as 'objective' measures to increase *accuracy* rather than efficiency. My analysis leaves no question about the falsity of such views. As we have seen, those who rely on the canons also often treat them as additional tipping reasons and advocate for using them so but not for using them in lieu of the balance of reasons (ie as exclusionary reasons). My analysis also debunks such approaches. The speculation as to whether the canons could be used in lieu of the balance of reasons and as efficiency-increasing devices is entirely mine and in light of my analysis of them as summary rules. It is a view that I have anticipated as a potential objection and have argued against—though, to my knowledge, it has no actual proponents.